

Geometrical organization of solutions to random linear Boolean equations

Thierry Mora and Marc Mézard

*Laboratoire de Physique Théorique et Modèles statistiques,
UMR 8626, CNRS and Université Paris Sud
Orsay Cedex, F-91405 (France)*

The random XORSAT problem deals with large random linear systems of Boolean variables. The difficulty of such problems is controlled by the ratio of number of equations to number of variables. It is known that in some range of values of this parameter, the space of solutions breaks into many disconnected clusters. Here we study precisely the corresponding geometrical organization. In particular, the distribution of distances between these clusters is computed by the cavity method. This allows to study the ‘ x -satisfiability’ threshold, the critical density of equations where there exist two solutions at a given distance.

Keywords: Message-passing algorithms, Typical-case computational complexity, Cavity and replica method.

I. INTRODUCTION

Constraint Satisfaction Networks (CSN) are problems involving many discrete variables, with values in a finite alphabet, related by low density constraints: each constraint involves a finite number of variables. This kind of problems arise in many branches of science, from statistical physics (spin or structural glasses [1]) to information theory (low-density parity-check (LDPC) codes [2, 3]) and combinatorial optimization (satisfiability, colouring [4]). The ‘thermodynamic limit’ of such problems is obtained when the number of variables and the number of constraints go to infinity, keeping their ratio, the density of constraints α , fixed. A lot of attention has been focused in recent years on the study of random CSN, both because of their practical interest in coding, and also as a means to study “typical case” complexity (as opposed to the traditional worst case complexity analysis). Many CSN are known to undergo a SAT-UNSAT phase transition when the density of constraints increases: there is a sharp threshold separating a SAT phase where all constraints can be satisfied with probability one in the thermodynamic limit from an UNSAT phase where, with probability one, there is no configuration of the variables satisfying all the constraints. While the existence of a sharp threshold has been proved by Friedgut [5] for satisfiability and colouring, there is no yet any rigorous proof of the widely accepted conjecture according to which the threshold density of constraints converges to a fixed value α_c in the thermodynamic limit.

Recent years have seen the upsurge of statistical physics methods in the study of CSN. In particular, the replica method and the cavity method have been used to study the phase diagram [6–8]. Their most spectacular results are some arguably exact (but not yet rigorously proved) expressions for α_c , and the existence of an intermediate SAT phase, in a region of constraint density $]\alpha_d, \alpha_c[$, where the space of solutions is split into many clusters, far away from each other. This clustering is an important building block of the theory: it is at the origin of the necessity to use the cavity method at the so-called one-step replica

symmetry breaking (1RSB) level; this method can be seen as a message-passing procedure and used as an algorithm for finding a SAT assignment of the variables. This algorithm, called survey propagation, turns out to be very powerful in satisfiability and colouring, and its effectiveness can be seen as one indirect piece of evidence in favour of clustering. On intuitive grounds, clustering is often held responsible for blocking many local search algorithms [9]. Although there does not exist any general discussion of this statement, this phenomenon was thoroughly investigated in the case of XORSAT [23].

The clustering effect can be studied in a more formal way by introducing the notion of x -satisfiability [10, 11]. A CSN with N variables is said x -satisfiable (x -SAT) if there exists a pair of SAT assignments of the variables which differ in a number of variables $\in [Nx - \epsilon(N), Nx + \epsilon(N)]$. Here x is the reduced distance, which we keep fixed as N goes to infinity. The resolution $\epsilon(N)$ has to be sub-linear in N : $\lim_{N \rightarrow \infty} \epsilon(N)/N = 0$, but its precise form is unimportant for our large N analysis. For example we can choose $\epsilon(N) = \sqrt{N}$. For many random CSN, it is reasonable to conjecture, in parallel with the existence of a satisfiability threshold, that x -satisfiability has a sharp threshold $\alpha_c(x)$ such that:

- if $\alpha < \alpha_c(x)$, a random formula is x -SAT almost surely.
- if $\alpha > \alpha_c(x)$, a random formula is x -UNSAT almost surely.

This conjecture has been proposed for k -satisfiability of random Boolean formulas where each clause involves exactly k variables with $k \geq 3$. So far only a weaker conjecture, analogous to Friedgut's theorem [5], has been established [11]. It states the existence of a non-uniform threshold $\alpha_c^{(N)}(x)$. Rigorous bounds on $\alpha_c(x)$ have been found in [11] for the k -satisfiability problem with $k \geq 8$, using moment methods developed in [12], but so far this x -satisfiability threshold has not been computed.

In this paper we compute the x -satisfiability threshold $\alpha_c(x)$ in the random XORSAT problem using the cavity method. This is a problem of random linear equations with Boolean algebra. It is important because many efficient error correcting codes are based on low-density parity-checks, the decoding of which involves precisely such linear systems. It is also one of the best understood case of CSN. In particular, efforts to extend the replica method [13] and the cavity method [14] to deal with models defined on finite-connectivity lattices, have resulted in the first exact (but non-rigorous) derivation of its phase diagram [15]. Later, a clear characterization of these clusters, combined with simple combinatoric arguments, gave a rigorous base to these predictions [16–18]. These works have computed the phase diagram in details and provide expressions for the two thresholds $\alpha_d < \alpha_c < 1$.

Our computation of $\alpha_c(x)$ confirms this known structure, and it also provides insight into the geometrical structure of clusters. We find that $\alpha_c(x)$ is non monotonic (see fig. 5), which confirms the existence of gaps in distances where there does not exist any pair of solutions.

The method used in our computation is in itself interesting. It turns out that it is not possible to compute $\alpha_c(x)$ directly, by fixing x and varying α . Instead, we work at a fixed value of α and introduce a probability distribution for pairs of SAT assignments, where the distance between the solutions plays the role of the energy. The computation of the entropy as a function of the energy, and more precisely the computation of the energies where it vanishes, then allows to reconstruct $\alpha_c(x)$. Our computation thus involves a mixture of hard constraints (the fact that the two assignments must satisfy the XORSAT formula), and soft constraints (the Boltzmann weight which depends on their distance). This is reflected in the structure of the cavity fields that solve this problem.

The remainder of this paper is organized as follows. The next section introduces some notations. In section III, we analyse classical Survey Propagation on XORSAT and show its equivalence with the “leaf removal” [18] or “decimation” [16] algorithm. This analysis allows to re-derive the phase diagram of XORSAT, and sets up useful notations and concepts for later computations. In section IV we perform a statistical mechanics analysis of weight properties in a single cluster using the cavity method. Section V applies this formalism to the computation of the cluster diameter, while section VI is devoted to the evaluation of inter-cluster distances. In section VII we sum up and discuss our results.

II. NOTATIONS AND DEFINITIONS

A XORSAT formula is defined on a string of N variables $x_1, x_2, \dots, x_N \in \{0, 1\}$ by a set of M parity checks of the form:

$$\sum_{i \in V(a)} x_i = y_a \pmod{2}, \quad \text{for all } a = 1, \dots, M \quad (1)$$

where $y_a \in \{0, 1\}$. Here $V(a) \subset \{1, \dots, N\}$ is the subset of variables involved in parity check a . Later on $i \in a$ shall be used as a shorthand for $i \in V(a)$.

Eq. (1) can be rewritten in the matricial form:

$$A\mathbf{x} = \mathbf{y} \pmod{2}, \quad A = \{A_{ia}\}_{i \in [N], a \in [M]} \quad (2)$$

where $A_{ia} = 1$ if $i \in a$ and $A_{ia} = 0$ otherwise. The pair $F = (A, \mathbf{y})$ defines the formula. Such a linear system can be solved in polynomial time by Gaussian elimination. If a formula has solutions, it is SAT; otherwise, it is UNSAT. The thermodynamics limit is $N \rightarrow \infty$, $M \rightarrow \infty$ with a fixed density of constraints $\alpha = M/N$.

In this paper we specialize to random k -XORSAT formulæ, where each equation involves a subset of k variables, chosen independently with uniform probability among the $\binom{N}{k}$ possible ones, and each y_a independently takes value 0 or 1 with probability 1/2. One important characterization of a XORSAT formula $F = (A, \mathbf{y})$ is the number $\mathcal{N}_N(F)$ of assignments of the Boolean variables \mathbf{x} which satisfy all the equations, and the corresponding entropy density

$$s_N(F) = \frac{1}{N} \log \mathcal{N}_N(F) \quad (3)$$

Logarithm is base 2 throughout the paper. Using a spin representation $\sigma_i = (-1)^{x_i}$, the k -XORSAT problem can also be mapped onto a spin glass model where interactions involve products of k spins (the variables $(-1)^{y_a}$ then play the role of quenched random exchange couplings) [15], and the question of whether a formula is SAT is equivalent to asking whether the corresponding spin-glass instance is frustrated.

Previous work [15–18] has shown that:

- for $\alpha < \alpha_d(k)$, the formula is SAT, almost surely (i.e. with probability $\rightarrow 1$ as $N \rightarrow \infty$). The solution set forms one big connected component, the entropy density concentrates at large N to $(N - M)/N = 1 - \alpha$; This phase is called the EASY-SAT phase.

- for $\alpha_d(k) < \alpha < \alpha_c(k)$, the formula is still SAT almost surely, but the solution set is made of an exponentially large (in N) number of components far away from each other (in the following we shall give a precise definition of these clusters); The entropy density also concentrates at large N to $(N - M)/N = 1 - \alpha$. This is the HARD-SAT phase.
- for $\alpha > \alpha_c(k)$ (with $\alpha_c(k) < 1$), the formula is UNSAT almost surely. The entropy is $-\infty$. This second transition is the usual SAT-UNSAT transition.

The fact that, throughout the SAT phase ($\alpha < \alpha_c(k)$), the entropy density concentrates to $1 - \alpha$ is not surprising: it can be understood as the fact that matrix A has rank M almost surely in the SAT phase. The intuitive reason is that, each time there exists a linearly dependent set of checks, the choice of y_a has probability 1/2 to lead to a contradiction. So the rank of A cannot differ much from M in the SAT phase. From the point of view of linear algebra, the existence of the clustered phase, i.e. the fact that the vector subspace of SAT assignments breaks into disconnected pieces, is more surprising, as is the discontinuity of $s_N(F)$ at the transition α_c . These two aspects are in fact related: the quantity which vanishes at the SAT-UNSAT transition is actually the log of the number of clusters of solutions, while each cluster keeps a finite volume.

We will study the geometric properties of the space of solutions for random k -XORSAT in the HARD-SAT phase using the notion of x -satisfiability. In terms of solutions of linear equations, we want to know if there exist two Boolean vectors \mathbf{x} and \mathbf{x}' which both satisfy $A\mathbf{x} = A\mathbf{x}' = \mathbf{y}$, where the Hamming distance $d_{\mathbf{x}, \mathbf{x}'} \equiv (\mathbf{x} - \mathbf{x}')^2 = Nx$. Clearly, if such a pair exists, $\mathbf{x} - \mathbf{x}'$ is solution to the homogeneous ('ferromagnetic') problem where $\mathbf{y} = \mathbf{0}$:

$$A(\mathbf{x} - \mathbf{x}') = \mathbf{0} \quad (4)$$

Therefore, a formula $F = (A, \mathbf{y})$ is x -SAT if and only if F is SAT and if there exists a solution \mathbf{x} to the homogeneous system $A\mathbf{x} = \mathbf{0}$ of weight $d_{\mathbf{x}, \mathbf{0}} \approx Nx$ (the *weight* is by definition the distance to $\mathbf{0}$). Note that for $x = 0$, this second condition is automatically fulfilled, and x -satisfiability is equivalent to satisfiability. This linear space structure also implies that the set of solutions looks the same seen from any solution in the SAT phase: the number of solutions at distance d of any given solution \mathbf{x}_0 is independent from \mathbf{x}_0 .

Distance properties can also be investigated directly by evaluating extremal distances between solutions. To that end we define three distances: (a) the cluster diameter d_1 , i.e. the largest Hamming distance between solutions belonging to the same cluster; this diameter is independent of the cluster; (b) the minimal and maximal inter-cluster distances d_2 and d_3 , i.e. the smallest (resp. largest) Hamming distance between solutions belonging to distinct clusters. All three distances are assumed to be self-averaging in the thermodynamic limit of the random problem: $x_1(\alpha) = d_1/N$, $x_2(\alpha) = d_2/N$ and $x_3(\alpha) = d_3/N$ shall denote the corresponding limits. In the particular case where k is even, the formula is invariant under the transformation $\mathbf{x} \leftrightarrow \mathbf{x} + \mathbf{1}$ (mod 2), which is reflected in terms of distances by a symmetry with respect to $x = 1/2$: $x \leftrightarrow 1 - x$. A direct consequence is that $x_3(\alpha) = 1 - x_2(\alpha)$, and that a fourth weight, defined as $1 - x_1(\alpha)$, will also come into play. These distance functions are related to the x -satisfiability threshold as follows: at fixed α , a formula is x -SAT almost surely iff

- $x \in [0, x_1(\alpha)] \cup [x_2(\alpha), x_3(\alpha)]$ when k is odd.
- $x \in [0, x_1(\alpha)] \cup [x_2(\alpha), 1 - x_2(\alpha)] \cup [1 - x_1(\alpha), 1]$ when k is even.

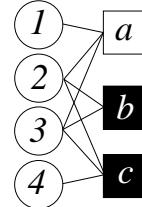
We will now compute x_1, x_2, x_3 with the cavity method.

III. LEAF REMOVAL AS AN INSTANCE OF SURVEY PROPAGATION

XORSAT formulæ are conveniently represented by factor graphs, called *Tanner graphs*, in which variables and checks form two distinct types of nodes, with the simple rule that the edge (i, a) between i and a is present if $i \in a$.

An example of a Tanner graph and its associated linear system is shown below:

- (a) $x_1 + x_2 + x_3 = 0 \pmod{2}$
- (b) $x_2 + x_3 = 1 \pmod{2}$
- (c) $x_2 + x_3 + x_4 = 1 \pmod{2}$



The number of variables involved in a check a , denoted by $|V(a)|$, is the degree of a in the factor graph. Here we study k -XORSAT where this degree is fixed to k . Similarly, if $V(i)$ denotes the set of parity checks in which i is represented, $|V(i)|$ is the degree of i in the factor graph. The degrees of checks are commonly referred to as *right-degrees*, and those of variables as *left-degrees*. The infinite-length (thermodynamic) limit is obtained by sending N and M to infinity while keeping the ratio $\alpha = M/N$ fixed. In this limit, the distribution of left-degrees is a Poisson law of parameter $k\alpha$: The probability of a variable having degree ℓ is $\pi_{k\alpha}(\ell)$, where $\pi_x(\ell) = \exp(-x)x^\ell/\ell!$.

Here we use the *leaf removal algorithm* (LR) in order to obtain a precise definition of the notion of “cluster” or “component” of solutions, one which is valid also for finite N . The algorithm proceeds as follows: pick a variable of degree one (called a *leaf*), remove it as well as the only check it is connected to. Continue the process until there remains no leaf. The interest of this algorithm is easily seen: a variable on a leaf can always be assigned in such a way that the (unique) check to which it is connected is satisfied.

The linear system remaining after leaf removal is independent of the order in which leaves are removed. It is called the *core*. A ‘core check’ is a check which only involves core variables. If the core is empty, the problem is trivially SAT. In general, given a solution of the core, one can easily reconstruct a solution of the complete formula by running leaf removal in the reverse direction, in a scheme which we refer to as *leaf reconstruction*. In this procedure, checks are added one by one along with their leaves, starting from the core. If an added check involves only one leaf, the value of that variable is determined uniquely so that the check is satisfied. If the number of leaves k' is greater than 1, one can choose the joint value of those leaves among $2^{k'-1}$ possibilities. The process is iterated until the complete factor graph has been rebuilt. Given a core solution, one can construct many solutions to the complete formula. Variables which are uniquely determined by the core solution are called *frozen*, and variables that can fluctuate are called *floppy*. Of course, by definition, the frozen part includes the core itself. A core solution defines a *cluster*. All solutions built from the same core solution belong to the same cluster. We shall see later how this definition fits in the intuitive picture that we sketched previously in terms of connectedness.

We propose here an alternative to the leaf removal algorithm, which also builds the core, but keeps actually more information. The approach is inspired by the cavity method, and is a special instance of Survey Propagation (SP) [7]. To each edge (i, a) , one assigns two numbers $\hat{m}_{a \rightarrow i}^t$ and $m_{i \rightarrow a}^t$ belonging to $\{0, 1\}$, updated as follows:

- At $t = 0$, $\hat{m}_{a \rightarrow i}^0 = 1$, $m_{i \rightarrow a}^0 = 1$ for all edges (i, a) .

- $m_{i \rightarrow a}^{t+1} = 1 - \prod_{b \in i-a} (1 - \hat{m}_{b \rightarrow i}^t)$.
- $\hat{m}_{a \rightarrow i}^t = \prod_{j \in a-i} m_{j \rightarrow a}^t$.
- Stop when $\hat{m}_{a \rightarrow i}^{t+1} = \hat{m}_{a \rightarrow i}^t$ for all (i, a) ,

Here $a \in i$ is a shorthand for $a \in V(i)$.

The interpretation of $m_{i \rightarrow a}^t = 1$ is: “variable i is constrained at time t in the absence of check a ”, and $\hat{m}_{a \rightarrow i}^t = 1$: “check a constrains variable i at time t ”. One also defines $M_i^t = 1 - \prod_{a \in i} (1 - \hat{m}_{a \rightarrow i}^t) \in \{0, 1\}$. This number indicates whether node i is constrained at time t ($M_i^t = 1$) or not ($M_i^t = 0$).

At $t = 0$, all variables are constrained. The algorithm consists in detecting the under-constrained variables, and propagating the information through the graph to simplify the formula. At the first step, only variables of degree one are affected: if i is of degree one and is connected to a , $m_{i \rightarrow a}^1 = 1 - \prod_{\emptyset} = 0$. This, in turn, gives freedom to a , which no longer constrains its other variables: $\hat{m}_{a \rightarrow j}^1 = 0$, for $j \in a-i$. This effectively removes a and i from the formula, just as in the leaf removal algorithm. In the subsequent steps of the iteration, will be considered as a *leaf* (in the LR sense), a variable i such that there exists exactly one $a \in i$ such that $\hat{m}_{a \rightarrow i}^t = 1$. In that case we have $m_{i \rightarrow a}^{t+1} = 0$, thus implementing a step of LR.

Let us add a word about the term “Survey Propagation” we have used so far. Analysis of the 1RSB cavity equations at zero temperature [18] (see [7] for a more complete discussion in the case of k -SAT) shows that cavity biases fall into two categories, depending on the edge we consider: either a warning is sent (compelling to take value 0 or 1 depending on the cluster, with probability one half for each), or no warning is sent. (In more technical terms, the survey propagation reduces to warning propagation). The first situation corresponds in our language to $\hat{m}_{a \rightarrow i} = 1$ and the second to $\hat{m}_{a \rightarrow i} = 0$. Similarly, we have $m_{i \rightarrow a} = 1$ if the cavity field is non-zero, and $m_{i \rightarrow a} = 0$ otherwise. Therefore our algorithm carries the same information as Survey Propagation.

The interest of SP over leaf removal is that it keeps track of the leaves which are uniquely determined by their check. For example, if two or more leaves are connected to the same check a at time t , at time $t+1$ one has $\hat{m}_{a \rightarrow i}^{t+1} = 0$ for all $i \in a$, reflecting the fact that a cannot uniquely determine the value of several leaves. Conversely, if a is connected to a unique leaf i and if one has: $m_{j \rightarrow a}^t = 1$ for all $j \in a-i$, then one gets $\hat{m}_{a \rightarrow i}^t = 1$, reflecting the fact that, the variables $\{x_j\}_{j \in a-i}$ being fixed in the absence of a , i is determined uniquely.

A little reasoning shows that when the algorithm stops ($t = t_f$), i is frozen iff $M_i^{t_f} = 1$, and i belongs to the core iff there exists at least two checks $a, b \in i$ such that $\hat{m}_{a \rightarrow i}^{t_f} = \hat{m}_{b \rightarrow i}^{t_f} = 1$. In the final state, we say that the directed edge $i \rightarrow a$ is frozen if $m_{i \rightarrow a} \equiv m_{i \rightarrow a}^{t_f} = 1$, and that $a \rightarrow i$ is frozen if $\hat{m}_{a \rightarrow i} \equiv \hat{m}_{a \rightarrow i}^{t_f} = 1$. In the opposite case, edges are called floppy. This version of SP is strictly equivalent to the *Belief Propagation* algorithm used for decoding Low-Density Parity-Check codes on the binary erasure channel, also called “Peeling decoder” in that context.

SP can be studied by density evolution in order to derive the phase diagram, as in [18]. Let us briefly survey this study for completeness. The statistics of messages at time t is described by two numbers:

$$v^t = \frac{1}{Mk} \sum_{(i,a)} \delta(m_{i \rightarrow a}^t, 0) \quad , \quad w^t = \frac{1}{Mk} \sum_{(i,a)} \delta(\hat{m}_{a \rightarrow i}^t, 0) , \quad (5)$$

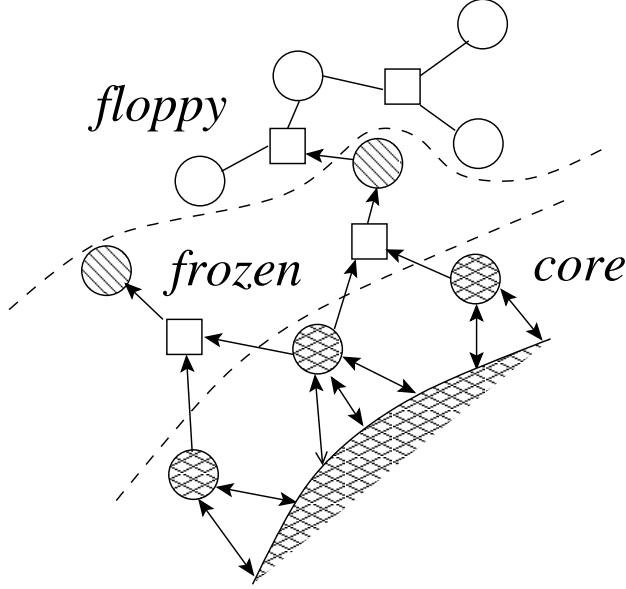


FIG. 1: A example of a fixed point of SP. Circles represent variable nodes, and squares check nodes. An arrow means that message m or \hat{m} has value 1, that is, that the directed edge is frozen when SP stops. Leaf removal propagates null messages from the outer leaves down to the core, while “leaf reconstruction” propagates non-null messages from the core up the frozen part.

where the sums run over all edges of the Tanner graph. When $N \rightarrow \infty$, these densities are governed by evolution equations:

$$\begin{aligned} v^{t+1} &= \sum_{\ell} \pi_{k\alpha}(\ell) (w^t)^{\ell} = \exp [-k\alpha(1 - w^t)] \\ w^t &= 1 - [1 - v^t]^{k-1}, \end{aligned} \quad (6)$$

which are initialized with $v^0 = w^0 = 0$. These equations are exact if the Tanner graph is a tree. In our case the graph is locally tree-like (it is a tree up to finite distance when seen from a generic point), and one could set up a rigorous proof of (6) using the methods developed in [19].

The fixed point of these equations is given by the *cavity equation*:

$$w = 1 - \{1 - e^{-k\alpha(1-w)}\}^{k-1}. \quad (7)$$

Setting $\lambda = k\alpha(1 - w)$, Eq. (7) can be rewritten as:

$$\lambda = k\alpha(1 - e^{-\lambda})^{k-1} \quad (8)$$

When $\alpha < \alpha_d$, the unique fixed point is $\lambda = 0$ (i.e. $w = 1$). This means that the core is empty. For $\alpha > \alpha_d$ however, there remains an extensive core of size

$$N_c = N \left[\sum_{\ell \geq 2} \pi_{k\alpha}(\ell) (1 - w^{\ell} - \ell w^{\ell-1}) \right] = N [1 - (1 + \lambda)e^{-\lambda}] \quad (9)$$

while the number of frozen variables is

$$N_f = N \left[\sum_{\ell \geq 2} \pi_{k\alpha}(\ell) (1 - w^\ell) \right] = N [1 - e^{-\lambda}] \quad (10)$$

The number of core checks is:

$$M_c = M(1 - v)^k = \alpha N [1 - e^{-\lambda}]^k. \quad (11)$$

The left-degree distribution (with respect to core checks) inside the core is given by a truncated Poissonian:

$$P_c(\ell) = \frac{1}{e^\lambda - 1 - \lambda} \frac{\lambda^\ell}{\ell!} \mathbb{I}(\ell \geq 2), \quad (12)$$

where \mathbb{I} is the indicator function.

One can show that the leaf removal algorithm conserves the uniformity of the ensemble. Therefore, the core formula is a random XORSAT formula with right-degree k and left-degree distribution $P_c(\ell)$ given by (12). The number of solution to such a formula is known to concentrate to its mean value when the size goes to infinity [17, 18]. In the case of the core formula, this number is simply $2^{N_c - M_c}$ if $N_c \geq M_c$, and 0 otherwise. Recalling that the complete formula has solutions if and only if the core formula does, we find that the SAT-UNSAT threshold α_c is given by the equation:

$$1 - (1 + \lambda)e^{-\lambda} = \alpha [1 - e^{-\lambda}]^k. \quad (13)$$

The number of clusters is characterized by the *complexity* or *configurational entropy*, that is the logarithm of the number of core solutions:

$$\Sigma(\alpha) = \frac{1}{N} \log(\# \text{ clusters}) = \frac{N_c - M_c}{N} = 1 - (1 + \lambda)e^{-\lambda} - \alpha [1 - e^{-\lambda}]^k \quad (14)$$

We recall that the group structure of the solution set implies that all clusters have the same internal structure. Their common internal entropy is therefore given by:

$$s_{\text{inter}} = 1 - \alpha - \Sigma(\alpha) \quad (15)$$

where we have used the fact that the total entropy is $1 - \alpha$.

Let us comment on the relationship between our definition of clusters and the more traditional one. Usually, clusters are defined as the “connected” components of the solution set, where connectedness is to be understood in the following way: two solutions are connected if one can go from one to the other by a sequence of solutions separated by a finite Hamming distance (when $N \rightarrow \infty$). To make contact with our own definition of clusters, one needs to prove two things. First, that two solutions built from the same core solution are connected. Second, that two core solutions are necessarily separated by an extensive Hamming distance ($\geq cN$, with c constant), which implies that two solutions built from two distinct core solutions are not connected. Both proves can be found in [18]. This reconciles our definition (which holds for any single instance of XORSAT) with the usual one (which only makes sense for infinite-length ensembles).

IV. DISTANCE LANDSCAPE: THERMODYNAMICAL APPROACH

As we have already observed, studying pairs of solutions is equivalent to studying solutions to the ferromagnetic problem. Indeed, if S denotes the affine subspace of solutions to $A\mathbf{x} = \mathbf{y}$, and S_0 the vector subspace of solutions to $A\mathbf{x} = \mathbf{0}$, we have:

$$S \times S = \{(\mathbf{x}', \mathbf{x}' + \mathbf{x}), (\mathbf{x}', \mathbf{x}) \in S \times S_0\} \quad (16)$$

In particular, distances in S are reflected by weights in S_0 . Therefore, in order to study the range of attainable distances between solutions, one just needs to study the range of possible weights in S_0 . To that end we set a thermodynamical framework in which the weight plays the role of an energy:

$$E(\mathbf{x}) \equiv |\mathbf{x}| = \sum_i \delta_{x_i, 1}. \quad (17)$$

The Boltzmann measure at temperature β^{-1} is thus defined by:

$$\mathbb{P}(\mathbf{x}, \beta) = \frac{1}{Z(\beta)} \prod_a \delta_{\mathbb{F}_2} \left(\sum_{i \in a} x_i, 0 \right) 2^{-\beta|\mathbf{x}|} \quad (18)$$

where the normalization constant $Z(\beta)$ is the partition function. The Dirac delta-function, here defined on the two-element field \mathbb{F}_2 , enforces that only configurations of S_0 are considered. Remarkably, this measure is formally similar to the one used to infer the most probable codeword under maximum-likelihood decoding in Low-Density Parity-Check (LDPC) codes on the Binary Symmetric Channel [20]. In fact, as we shall see soon, some of the methods used to solve both problems share common aspects.

A very useful scheme for estimating marginal probabilities in models defined on sparse graphs is the cavity method [14], which we have already mentioned in the previous section. Let $p_{i \rightarrow a}^x$ be the probability that $x_i = x$ under the measure defined by (18), where the link (i, a) has been removed. The replica symmetric (RS) cavity method consists in computing the cavity marginals $p_{i \rightarrow a}^x$ (viewed as variable-to-check messages) using a closed set of equations where check-to-variable messages are also introduced as intermediate quantities. These second-kind messages are denoted by $q_{a \rightarrow i}^x$ and are proportional to the probability that $x_i = x$ when i is connected to a only. Messages are updated until convergence with the following rules:

$$p_{i \rightarrow a}^{x_i} = \frac{1}{Z_{i \rightarrow a}} \prod_{b \in i - a} q_{b \rightarrow i}^{x_i} 2^{-\beta \delta_{x_i, 1}} \quad (19)$$

$$q_{a \rightarrow i}^{x_i} = \sum_{\{x_j\}_{j \in a - i}} \prod_{j \in a - i} p_{j \rightarrow a}^{x_j} \delta_{\mathbb{F}_2} \left(\sum_{j \in a} x_j, 0 \right) \quad (20)$$

where $Z_{i \rightarrow a}$ is a normalization constant. When convergence is reached, marginal probabilities are obtained as:

$$p_i^{x_i} \equiv \sum_{\{x_j\}_{j \neq i}} \mathbb{P}(\mathbf{x}, \beta) = \frac{1}{Z_{i+a \in i}} \prod_{a \in i} q_{a \rightarrow i}^{x_i} 2^{-\beta \delta_{x_i, 1}} \quad (21)$$

where $Z_{i+a \in i}$ is also a normalization constant. Continuing the analogy with codes, it is interesting to note that these cavity equations are identical [21] to the Belief Propagation

(BP) equations [22] used to decode messages with LDPC codes on the Binary Symmetric Channel.

It turns out that cavity equations (19), (20) do not admit a unique solution, as one would expect if the system were replica symmetric. Instead, let us show that they admit exactly one solution for each cluster. In a given cluster denoted by \mathbf{c} , let us denote by c_i the value of a frozen variable i . There exists a solution to (19), (20), where, for every frozen variable i :

$$\begin{aligned} p_{i \rightarrow a}^x &= \delta_{x, c_i} && \text{if } i \rightarrow a \text{ frozen} \\ q_{a \rightarrow i}^x &= \delta_{x, c_i} && \text{if } a \rightarrow i \text{ frozen,} \end{aligned} \quad (22)$$

In order to show that this is a solution, let us use the SP messages, which provide information on how the fixing of the core solution forces the values of frozen variables. For example $m_{i \rightarrow a} = 1$ indicates that x_i is entirely determined by the core solution, supposing that the edge (i, a) has been removed. Consider the SP fixed point relations

$$\begin{aligned} \hat{m}_{a \rightarrow i} &= \prod_{j \in a-i} m_{j \rightarrow a}, \\ m_{i \rightarrow a} &= 1 - \prod_{b \in i-a} (1 - \hat{m}_{b \rightarrow i}). \end{aligned} \quad (23)$$

They are in fact *contained* in the cavity equations Eqs. (19), (20). In fact, the iteration of cavity equations allows to identify the frozen edges, irrespectively of the cluster the system falls into.

But the cavity equations also contain ‘fluctuating’ messages, where p^x and q^x are in $]0, 1[$, which are *de facto* restricted to the floppy part. We parametrize them by the cavity fields and biases:

$$\beta h_{i \rightarrow a}^{\mathbf{c}} = \log \frac{p_{i \rightarrow a}^0}{p_{i \rightarrow a}^1}, \quad \beta u_{a \rightarrow i}^{\mathbf{c}} = \log \frac{q_{a \rightarrow i}^0}{q_{a \rightarrow i}^1} \quad (24)$$

which satisfy the equations:

$$h_{i \rightarrow a}^{\mathbf{c}} = \sum_{b \in i-a} u_{b \rightarrow i}^{\mathbf{c}} + 1 \quad \text{with } i \rightarrow a \text{ floppy,} \quad (25)$$

$$\beta u_{a \rightarrow i}^{\mathbf{c}} = 2 \operatorname{atanh} \left[\prod_{j \in a^n - i} \tanh(\beta h_{j \rightarrow a}^{\mathbf{c}} / 2) \prod_{j \in a^f - i} (-1)^{c_j} \right] \quad \text{with } a \rightarrow i \text{ floppy} \quad (26)$$

Note that cavity messages $h_{i \rightarrow a}^{\mathbf{c}}$ and $u_{a \rightarrow i}^{\mathbf{c}}$ now depend explicitly on the considered cluster, and are uniquely determined by it.

The multiplicity of solutions to RS cavity equations is a clear sign that the replica symmetry is broken. The main lesson from this discussion is that solutions can fluctuate according to two hierarchical levels of statistics: the first level deals with fluctuations inside a single cluster, i.e. fluctuations on the floppy part, while the second level deals with the choice of the cluster. The reduced cavity equations (25), (26) correctly describe the first level¹, when

¹ Although the RS Ansatz is unable to describe the whole system, it can reasonably be assumed to be valid on a single cluster.

the system is forced to live in cluster \mathbf{c} . This leads to defining a new probability measure and partition function, restricted to \mathbf{c} :

$$Z_{\mathbf{c}}(\beta) = \sum_{\mathbf{x} \in \mathbf{c}} 2^{-\beta \sum_{i=1}^N \delta_{x_i,1}} \quad (27)$$

By construction, this system is characterized by the fixing of the frozen edges (22) and by the reduced cavity equations (25), (26). The second level of statistics, i.e. the statistics over the clusters, is appropriately handled by an 1RSB calculation, and will be the object of section VI. We first focus on the properties of single clusters under the measure defined by (27).

The cavity method comes with a technique to estimate the log of the partition functions, also called potential in our case:

$$\phi(\beta) = -\frac{1}{N} \log Z(\beta) \quad (28)$$

(Note that this quantity differs from the usual free energy by a factor β). It can be computed within the RS Ansatz by the Bethe formula [21]:

$$N\phi(\beta) = \sum_i \Delta\phi_{i+a \in i} - (k-1) \sum_a \Delta\phi_a \quad (29)$$

where

$$\begin{aligned} \Delta\phi_{i+a \in i} &= -\log Z_{i+a \in i} = -\log \sum_{x_i} \prod_{a \in i} q_{a \rightarrow i}^x 2^{-\beta \delta_{x_i,1}} \\ \Delta\phi_a &= -\log \sum_{\{x_i\}_{i \in a}} \prod_{i \in a} p_{i \rightarrow a}^{x_i} \delta_{\mathbb{F}_2} \left(\sum_{i \in a} x_i, 0 \right) \end{aligned} \quad (30)$$

This formula has a rather simple interpretation: $\Delta\phi_{i+a \in i}$ is the contribution of i and its adjacent checks to the potential. When these contributions are summed, each check is counted k times, whence the need to subtract $k-1$ times the contribution of each check $\Delta\phi_a$. Also note that this expression is variational: it is stationary in the messages $\{p_{i \rightarrow a}\}$ as soon as the cavity equations (19), (20) are satisfied.

The RS Ansatz is valid in a single cluster. The single cluster potential $\phi_{\mathbf{c}}(\beta) = -\frac{1}{N} \log Z_{\mathbf{c}}(\beta)$ can therefore be computed by plugging Eqs. (22), (25) and (26) into the Bethe formula (30), provided one uses the messages corresponding to one given cluster \mathbf{c} . When one is restricted to a single cluster \mathbf{c} , the range of possible weights is $[x_{\mathbf{c}}, X_{\mathbf{c}}]$. The minimal and maximal weights can be obtained by sending $\beta \rightarrow \pm\infty$. For $\beta \rightarrow \infty$, the second cavity equation (26) simplifies to:

$$u_{a \rightarrow i}^{\mathbf{c}} = \mathcal{S} \left(\prod_{j \in a^{nf}-i} h_{j \rightarrow a}^{\mathbf{c}} \prod_{j \in a^f-i} (-1)^{e_j} \right) \min_{j \in a^{nf}-i} |h_{j \rightarrow a}^{\mathbf{c}}| \quad \text{with } a \rightarrow i \text{ floppy} \quad (31)$$

where $\mathcal{S}(x) = 1$ if $x > 0$, -1 if $x < 0$ and 0 if $x = 0$.

The “ground state energy”, i.e. the minimal weight in \mathbf{c} , is obtained as:

$$x_{\mathbf{c}} = \lim_{\beta \rightarrow \infty} \partial_{\beta} \phi_{\mathbf{c}}(\beta) = \frac{1}{N} \sum_{i \text{ floppy}}^N \frac{1 - \mathcal{S}(\sum_{a \in i} u_{a \rightarrow i}^{\mathbf{c}} + 1)}{2} + \frac{1}{N} \sum_{i \text{ frozen}} \delta_{c_i,1} \quad (32)$$

The $\beta \rightarrow -\infty$ limit yields very similar equations. These equations will be analyzed in the next section.

Let us also write down the equations giving the potential, which will be used in sect.VI.

$$N\phi_{\mathbf{c}}(\beta) = \sum_i \Delta\phi_{i+a \in i}^{\mathbf{c}} - (k-1) \sum_a \Delta\phi_a^{\mathbf{c}} \quad (33)$$

$$\lim_{\beta \rightarrow -\infty} \frac{1}{\beta} \Delta\phi_{i+a \in i}^{\mathbf{c}} \equiv \Delta x_{i+a \in i}^{\mathbf{c}}, \quad \lim_{\beta \rightarrow -\infty} \frac{1}{\beta} \Delta\phi_a^{\mathbf{c}} \equiv \Delta x_a^{\mathbf{c}} \quad \text{with} \quad (34)$$

$$\Delta x_{i+a \in i}^{\mathbf{c}} = \frac{1}{2} \left(\sum_{a \in i} |u_{a \rightarrow i}^{\mathbf{c}}| + 1 - \left| \sum_{a \in i} u_{a \rightarrow i}^{\mathbf{c}} + 1 \right| \right) \quad \text{if } i \text{ is floppy} \quad (35)$$

$$\Delta x_{i+a \in i}^{\mathbf{c}} = \sum_{a \in i^{nf}} |u_{a \rightarrow i}^{\mathbf{c}}| \vartheta(-u_{a \rightarrow i}^{\mathbf{c}}) \quad \text{if } i \text{ is frozen and } c_i = 0 \quad (36)$$

$$\Delta x_{i+a \in i}^{\mathbf{c}} = 1 + \sum_{a \in i^{nf}} |u_{a \rightarrow i}^{\mathbf{c}}| \vartheta(u_{a \rightarrow i}^{\mathbf{c}}) \quad \text{if } i \text{ is frozen and } c_i = 1 \quad (37)$$

$$\Delta x_a^{\mathbf{c}} = \vartheta \left(- \prod_{i \in a^{nf}} h_{i \rightarrow a}^{\mathbf{c}} \prod_{i \in a^f} (-1)^{c_i} \right) \min_{i \in a^{nf}} |h_{i \rightarrow a}^{\mathbf{c}}| \quad (38)$$

V. DIAMETER

With our formalism, computing the cluster diameter boils down to computing the maximal weight in cluster **0** (the cluster containing **0**). The relevant partition function for this task is:

$$Z_{\mathbf{0}}(\beta) = 2^{-N\phi_{\mathbf{0}}(\beta)} = \sum_{\mathbf{x} \in \mathbf{0}} \delta_{\mathbb{F}_2} \left(\sum_{i \in a} x_i, 0 \right) 2^{-\beta \sum_{i=1}^N \delta_{x_i, 1}} \quad (39)$$

When $\beta \rightarrow -\infty$, the solution of the cavity equations corresponding to cluster **0** is characterized by:

$$\begin{aligned} p_{i \rightarrow a}^x &= \delta_{x,0} \quad \text{if } i \rightarrow a \text{ frozen,} \\ q_{a \rightarrow i}^x &= \delta_{x,0} \quad \text{if } a \rightarrow i \text{ frozen,} \\ h_{i \rightarrow a} &= \sum_{b \in i-a} u_{b \rightarrow i} + 1 \quad \text{if } i \rightarrow a \text{ floppy,} \\ u_{a \rightarrow i} &= -\mathcal{S} \left[\prod_{j \in a^{nf}-i} (-h_{j \rightarrow a}) \right] \min_{j \in a^{nf}-i} |h_{j \rightarrow a}| \quad \text{if } a \rightarrow i \text{ floppy} \end{aligned} \quad (40)$$

and the maximum weight d_1 is given by:

$$d_1 = \lim_{\beta \rightarrow -\infty} \partial_{\beta} \phi_{\mathbf{0}}(\beta) = \sum_{i \text{ floppy}}^N \frac{1 + \mathcal{S} (\sum_{a \in i} u_{a \rightarrow i} + 1)}{2} \quad (41)$$

These equations are presented for single XORSAT formulæ, and can be solved by simple iteration of the corresponding message-passing rules. In practice however, in the regime

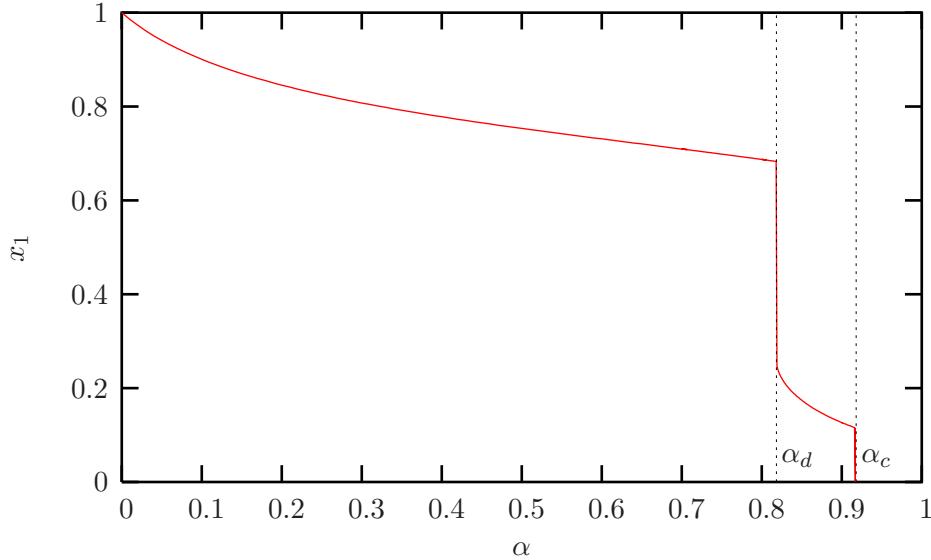


FIG. 2: Diameter of a cluster of solutions. When one decreases α below α_d all clusters aggregate into one big cluster, thus explaining the discontinuity.

where α is near (but smaller than) α_d , one does not always reach convergence. This is arguably due to the hard nature of XORSAT constraints, as it was pointed out in [23]: as one nears the dynamical transition, hopping from one solution to the other requires an increasing (yet sub-extensive) number of changes, making the sampling of solutions difficult. To circumvent this problem, we can work directly in the infinite-length limit by considering the probability distribution functions (pdfs) of each kind of message:

$$\begin{aligned} P(h) &= \frac{1}{Mk} \sum_{(i,a)} \delta_{h,h_i \rightarrow a} \\ Q(u) &= \frac{1}{Mk} \sum_{(i,a)} \delta_{u,u_a \rightarrow i} \end{aligned} \quad (42)$$

When $N \rightarrow \infty$, self-consistency equations for these distributions read:

$$\begin{aligned} P(h) &= \sum_{\ell} \pi_{k\alpha w}(\ell) \int \prod_{a=1}^{\ell} du_a Q(u_a) \delta \left(h - \sum_{a=1}^{\ell} u_a - 1 \right) \\ Q(u) &= \frac{1}{w} \sum_{i=1}^{k-1} \binom{k-1}{i} v^i (1-v)^{k-1-i} \int \prod_{j=1}^i dh_j P(h_j) \delta \left[u + \mathcal{S} \left(\prod_{j=1}^i (-h_j) \right) \min_j |h_j| \right] \end{aligned} \quad (43)$$

and one has:

$$x_1(\alpha) = \lim_{N \rightarrow \infty} \frac{d_1}{N} = e^{-\lambda} \int dh P(h) \frac{1 + \mathcal{S}(h)}{2} \quad (44)$$

These equations can be solved with a population dynamics algorithm [14]. In Fig. (2), we represent the maximal diameter x_1 as a function of α .

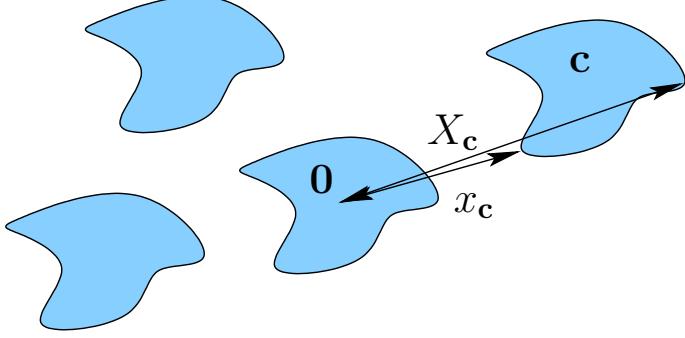


FIG. 3: Pictorial representation of the clustered space of solutions around $\mathbf{0}$ in the N -dimensional hypercube. For a cluster \mathbf{c} , the minimal and maximal distances $x_{\mathbf{c}}$ and $X_{\mathbf{c}}$ are depicted.

VI. MINIMAL AND MAXIMAL DISTANCES BETWEEN CLUSTERS

In section IV we have set up the formalism for computing the minimal and the maximal weights in a given cluster \mathbf{c} using the cavity method. In order to evaluate the minimal and maximal weights in *all* clusters except $\mathbf{0}$, we resort to a statistical treatment of the cavity equations. This scheme is known as the 1RSB cavity method in the replica language. We first specialize to the case of minimal weights, the other case being formally equivalent. We already know that the number of clusters grows exponentially with N . Here we further assume that the number of clusters with a given minimal weight $x_{\mathbf{c}}$ is exponential in N , and we define the complexity

$$\sum_{\mathbf{c} \neq \mathbf{0}} \delta(x, x_{\mathbf{c}}) = 2^{N\Sigma_m(x)}. \quad (45)$$

To this quantity we associate the 1RSB potential

$$2^{N\psi_m(y)} = \sum_{\mathbf{c} \neq \mathbf{0}} 2^{-Ny x_{\mathbf{c}}} = \int dx 2^{N(\Sigma_m(x) - yx)}. \quad (46)$$

When N is large, a saddle-point evaluation of this quantity yields:

$$\psi_m(y) = \min_x [yx - \Sigma_m(x)] = yx^* - \Sigma_m(x^*) \quad \text{with} \quad y = \partial_x \Sigma_m(x^*) \quad (47)$$

$\psi_m(y)$ is thus related to $\Sigma_m(x)$ by a Legendre transformation. In terms of statistical mechanics, m is an inverse temperature coupled to the “energy” $x_{\mathbf{c}}$; the complexity plays the role of a micro-canonical entropy, and the potential is equivalent to a free energy, up to a factor m . The minimal weight in all clusters (expect $\mathbf{0}$) is given by the smallest x such that $\Sigma_m(x) \geq 0$. Our goal is now to compute $\psi_m(y)$, and to infer $\Sigma_m(x)$ by inverse Legendre transformation.

We proceed to the statistical analysis of the cavity equations under Boltzmann measure $2^{-Ny x_{\mathbf{c}}}$. This amounts to writing 1RSB cavity equations, where messages are distributions of RS messages over all clusters. The distribution of messages on floppy edges is described by the two pdfs:

$$P^{i \rightarrow a}(h) = \langle \delta(h, h_{i \rightarrow a}^{\mathbf{c}}) \rangle \quad (48)$$

$$Q^{a \rightarrow i}(u) = \langle \delta(u, u_{a \rightarrow i}^{\mathbf{c}}) \rangle. \quad (49)$$

The average $\langle \cdot \rangle$ is performed with the aforementioned measure on clusters, with the implicit assumption that the edge (i, a) has been removed. On frozen edges, messages are trivial, but their values depend on the considered cluster. We thus define for frozen edges:

$$P_0^{i \rightarrow a} = \langle \delta(p_{i \rightarrow a}^0, 1) \rangle \quad P_1^{i \rightarrow a} = 1 - P_0^{i \rightarrow a} \quad (50)$$

$$Q_0^{a \rightarrow i} = \langle \delta(q_{a \rightarrow i}^0, 1) \rangle \quad Q_1^{a \rightarrow i} = 1 - Q_0^{a \rightarrow i} \quad (51)$$

In order to write a closed set of equations for these probability distributions, we need to know how the Boltzmann weight $2^{-Ny_{xc}}$ biases the message-passing procedure: when a field $h_{i \rightarrow a}$ is estimated as a function of its “grand-parents” ($\{h_{j \rightarrow b}\}$, $j \in b - i$, $b \in i - a$), a re-weighting term $2^{-y\Delta x_{i \rightarrow a}}$ is associated to it [7, 14], where $\Delta x_{i \rightarrow a}$ is the contribution of i and its adjacent checks (except a) to the total weight. This contribution is obtained as $\Delta x_{i+a \in i}$ in Eq. (35)-(37), but with a removed.

The 1RSB cavity equations read:

- $i \rightarrow a$ frozen:

$$\begin{aligned} P_0^{i \rightarrow a} &= \frac{1}{Z_{i \rightarrow a}} \prod_{b \in i^f - a} Q_0^{b \rightarrow i} \int \prod_{b \in i^f - a} du_{b \rightarrow i} Q^{b \rightarrow i}(u_{b \rightarrow i}) 2^{-y \sum_{b \in i^f - a} |u_{b \rightarrow i}| \vartheta(-u_{b \rightarrow i})} \\ P_1^{i \rightarrow a} &= \frac{1}{Z_{i \rightarrow a}} \prod_{b \in i^f - a} Q_1^{b \rightarrow i} \int \prod_{b \in i^f - a} du_{b \rightarrow i} Q^{b \rightarrow i}(u_{b \rightarrow i}) 2^{-y(1 + \sum_{b \in i^f - a} |u_{b \rightarrow i}| \vartheta(u_{b \rightarrow i}))} \end{aligned} \quad (52)$$

- $i \rightarrow a$ floppy:

$$\begin{aligned} P^{i \rightarrow a}(h) &= \frac{1}{Z_{i \rightarrow a}} \int \prod_{b \in i - a} du_{b \rightarrow i} Q^{b \rightarrow i}(u_{b \rightarrow i}) 2^{-y/2(\sum_{b \in i - a} |u_{b \rightarrow i}| + 1 - |\sum_{b \in i - a} u_{b \rightarrow i} + 1|)} \\ &\times \delta\left(h - 1 - \sum_{b \in i - a} u_{b \rightarrow i}\right) \end{aligned} \quad (53)$$

(here and in the previous equations $Z_{i \rightarrow a}$ is a normalization constant)

- $a \rightarrow i$ frozen:

$$Q_0^{a \rightarrow i} = \frac{1 + \prod_{j \in a - i} (2P_0^{j \rightarrow a} - 1)}{2} \quad (54)$$

- $a \rightarrow i$ floppy:

$$\begin{aligned} Q^{a \rightarrow i}(u) &= \sum_{\{c_j=0,1\}} \prod_{\substack{j \in a^f - i \\ j \in a^f - i}} P_{c_j}^{j \rightarrow a} \int \prod_{j \in a^f - i} dh_{j \rightarrow a} P^{j \rightarrow a}(h_{j \rightarrow a}) \\ &\times \delta\left[u - \mathcal{S}\left(\prod_{j \in a^f - i} h_{j \rightarrow a} \prod_{j \in a^f - i} (-1)^{c_j}\right) \min_{j \in a^f - i} |h_{j \rightarrow a}|\right] \end{aligned} \quad (55)$$

The potential $\psi_m(y)$ is obtained by a Bethe-like formula [7]:

$$N\psi_m(y) = \sum_i \Delta\psi_{i+a \in i} - (k-1) \sum_a \Delta\psi_a \quad (56)$$

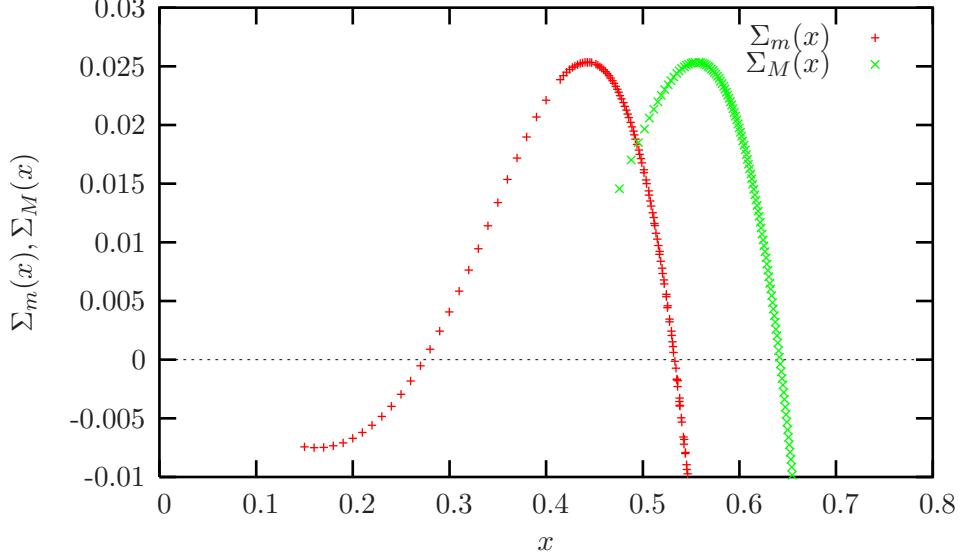


FIG. 4: Minimal and maximal distance complexities as a function of the reduced distance x , for $k = 3$, $N = 10000$ and $M = 8600$.

with

$$\begin{aligned}
 \Delta\psi_{i+a \in i} &= -\log \langle 2^{-y\Delta x_{i+a \in i}} \rangle = -\log \mathcal{Z}_{i+a \in i} \\
 \Delta\psi_a &= -\log \langle 2^{-y\Delta x_a} \rangle \\
 &= -\log \frac{1 + \prod_{i \in a} (2P_0^{i \rightarrow a} - 1)}{2} \quad \text{if } a \in \text{core} \\
 &= -\log \sum_{\substack{\{c_i=0,1\} \\ i \in a^f}} \prod_{j \in a^f} P_{c_i}^{i \rightarrow a} \int \prod_{i \in a^{nf}} dh_{i \rightarrow a} P^{i \rightarrow a}(h_{i \rightarrow a}) \\
 &\quad \times \exp \left[-y \log(2) \vartheta \left(- \prod_{i \in a^{nf}} h_{i \rightarrow a} \prod_{i \in a^f} (-1)^{c_i} \right) \min_{i \in a^{nf}} |h_{i \rightarrow a}| \right] \quad \text{otherwise}
 \end{aligned} \tag{57}$$

where $\mathcal{Z}_{i+a \in i}$ is defined as $\mathcal{Z}_{i \rightarrow a}$ but in the presence of a .

Like in the diameter calculation, 1RSB cavity equations can be interpreted as message-passing update rules, with the difference that messages are now surveys over all clusters. The output of that procedure is the minimal distance complexity $\Sigma_m(x)$, obtained as the inverse Legendre transform of $\psi_m(y)$. We refer to the corresponding algorithm as “distance survey propagation”. The same procedure can be implemented in the $\beta \rightarrow -\infty$ limit, and yields the maximal distance complexity:

$$\Sigma_M(x) = \frac{1}{N} \log \sum_{\mathbf{c} \neq 0} \delta(x, X_{\mathbf{c}}), \tag{58}$$

where $X_{\mathbf{c}}$ is the maximal weight in cluster \mathbf{c} (see Fig. 3). Note that in the particular case where $y = 0$, which corresponds to a uniform measure over the clusters, classical SP is recovered for both versions of the algorithm (minimal and maximal distance): in that limit

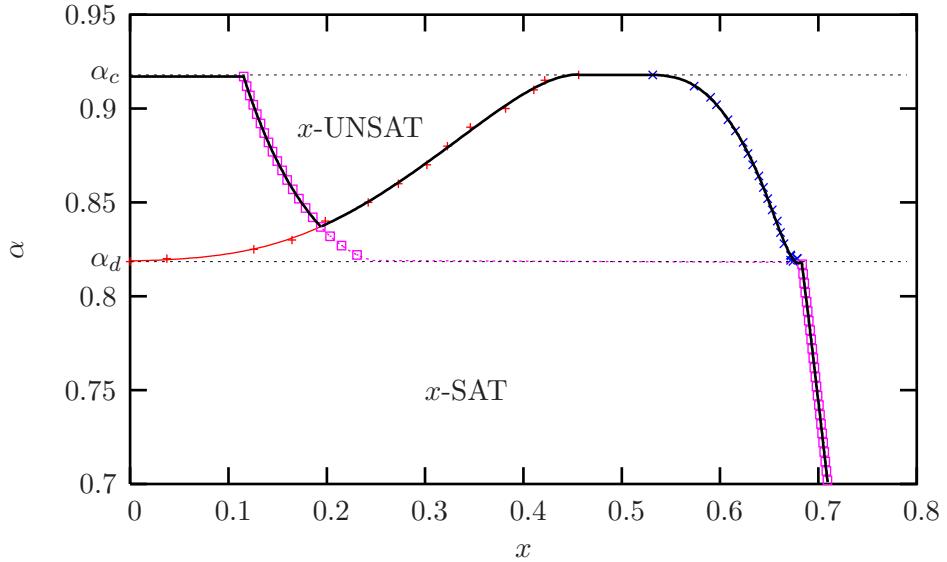


FIG. 5: Phase diagram of the 3-XORSAT problem in the (x, α) plane. The cluster diameter (\square), as well as minimal (+) and maximal (\times) distances between solutions of distinct clusters, are represented. The thick line is the x -satisfiability threshold.

we have $Q_0^{a \rightarrow i} = P_0^{i \rightarrow a} = 1/2$, and the calculation of $\psi_m(0)$ and $\psi_M(0)$ gives back $-\Sigma(\alpha)$, the total complexity (14), as expected.

The practical implementation of distance-SP demands particular care when small distances are considered: it turns out that distance complexities $\Sigma_m(x)$ and $\Sigma_M(x)$ are not concave, which entails that the functions $\psi_m(y)$ and $\psi_M(y)$ are multivalued in a certain range of y . A way to circumvent this problem (already used in [24]) is to keep the weight $x = \partial_y \psi_m(y)$ fixed after each iteration, and to deduce y accordingly. Here is how the algorithm proceeds for a given reduced weight x :

1. Run classical SP.
2. Initialize all floppy and frozen messages $\{P_{i \rightarrow a}\}$, $\{Q_{a \rightarrow i}\}$ to random values. Choose a (reasonable) value for y .
3. Until convergence is reached, do:
 - Update all $a \rightarrow i$ messages $\{Q_{a \rightarrow i}\}$, and then all $i \rightarrow a$ messages $\{P_{i \rightarrow a}\}$ at inverse temperature y .
 - Find y such that $x = \partial_y \psi_m(y, \{P_{i \rightarrow a}\}, \{Q_{a \rightarrow i}\})$ by the secant method, $\{P_{i \rightarrow a}\}$ and $\{Q_{a \rightarrow i}\}$ being fixed.
4. Compute $\psi_m(y, \{P_{i \rightarrow a}\}, \{Q_{a \rightarrow i}\})$ as well as its derivative and deduce $\Sigma_m(x) = yx - \psi_m(y)$.

Note that since the messages are pdfs themselves, the update of each of them in step 3 is performed by a population dynamics sub-routine.

Fig. 4 shows the minimal and maximal weight complexities $\Sigma_m(x)$ and $\Sigma_M(x)$ for a random 3-XORSAT formula with $N = 10000$ and $M = 8600$. These complexities can

be regarded as kinds of weight enumerator functions for clusters. Their fluctuations from formula to formula can be significant (15%), even for large system sizes ($N = 10000$).

An average version (density evolution) of distance-SP can also be implemented for random k -XORSAT, in the same spirit as Eq. (43). Such a computation involves distributions (on edges) of distributions (on clusters), and can be solved by population dynamics, where each element of the population is itself a population. The zeros of $\overline{\Sigma_m(x)}$ and $\overline{\Sigma_M(x)}$ thus obtained yield the minimal and maximal inter-cluster distances $x_2(\alpha)$ and $x_3(\alpha)$, respectively, as shown in Fig. 5. Together with the cluster diameter $x_1(\alpha)$ computed in section V, these values are used to construct the x -satisfiability threshold.

Our algorithm can in principle be run on any system of Boolean linear equations, and is expected to give reasonable results provided that the loops of the underlying Tanner graph are large. The case of LDPC codes is of particular interest because it allows several simplifications and has been extensively studied from both the combinatorial [25] and statistical physics [24, 26] point of view. LDPC codes are homogeneous Boolean linear systems where parity checks and variables may have arbitrary degree distributions, with the restriction that variables should always have degrees no less than 2. This implies that the leaf removal algorithm is inefficient on such linear systems: all variables belong to the core, and are frozen. In particular, each cluster is made of one unique solution: the cluster diameter is 0, and the minimal and maximal inter-cluster distances coincide. Their common complexity $\Sigma_m(x) = \Sigma_M(x)$ is often called ‘weight enumerator exponent’ and is an important property of ensembles of codes. Translated into our formalism, this means that all messages are frozen, and the distance-SP algorithm simplifies dramatically:

$$P_0^{i \rightarrow a} = \frac{1}{\mathcal{Z}_{i \rightarrow a}} \prod_{b \in i^f - a} Q_0^{b \rightarrow i}, \quad P_1^{i \rightarrow a} = \frac{1}{\mathcal{Z}_{i \rightarrow a}} \prod_{b \in i^f - a} Q_1^{b \rightarrow i} 2^{-y} \quad (59)$$

$$Q_0^{a \rightarrow i} = \frac{1 + \prod_{j \in a - i} (2P_0^{j \rightarrow a} - 1)}{2} \quad (60)$$

Not surprisingly, the density evolution analysis of this simplified algorithm yields the same equations as those obtained with the replica method in [24, 26].

VII. CONCLUSION AND DISCUSSION

We have applied the cavity method to estimate extremal distances between solutions of random linear systems with large girth in the clustered phase. Our results are used to compute the x -satisfiability threshold of the random k -XORSAT problem. The notion of x -satisfiability, which tells us whether one can find a pair of solutions separated by a Hamming distance x , was introduced in the context of another constraint satisfaction problem, k -SAT, where it was used to give rigorous evidence in favor of the clustering phenomenon [10].

Although k -XORSAT is a rather simple problem, it displays a very similar phase diagram as harder problems such as k -SAT or q -colorability. In particular, its clustered phase is well defined and understood. That said, finding extremal distances in the solution space of linear Boolean equations is a hard task in general: for instance, the decision problem associated with finding the minimal weight of LDPC codes is NP-complete [27].

We were able to compute three quantities: the cluster diameter, as well as the minimal and maximal inter-cluster distances. We believe our method to give a good approximation for

systems with large girth, and to be exact in the thermodynamic limit for random XORSAT. In the line of Survey Propagation, we devised a series of algorithms for these tasks, which explicitly exploit the clustered structure of the solution space. More precisely, the space of solutions is characterized by two hierarchical levels of fluctuations: inside and between clusters. In k -XORSAT, these two kinds of fluctuations are carried by two disjoint sets of variables, and our algorithms explicitly distinguish between these two types of variables. In the special case of LDPC codes, the point-like nature of clusters much simplifies the equations, and previous expressions of the weight enumerator exponent obtained by the replica method are recovered.

The method presented here offers a number of generalizations. In particular, it could be used at finite temperature to yield the full weight enumerator function. More interestingly, it could be adapted to deal with other CSN, such as k -SAT, for which only bounds are known; unfortunately, numerical computations are in that case much heavier, albeit formally similar. Let us mention that a similar approach was followed in [28] in the case of q -colorability, with the difference that distances were estimated from a reference configuration (which is not a solution) instead of considering distances between solutions.

Our work studies the geometrical properties of the solution space by taking explicitly into account fluctuations inside clusters, captured by the ‘evanescent fields’. This very general approach, already explored in [28], allows to gain a better understanding of the fine structure of the clustered phase, and seems to us a promising direction for future work. Also, with similar tools, decimation schemes such as the one introduced in [7] could be used to select solutions or clusters with particular properties.

We would like to thank Andrea Montanari for sharing the numerical trick used in the replica evaluation of the weight enumerator function of LDPC codes [24]. This work has been supported in part by the EU through the network MTR 2002-00319 ‘STIPCO’ and the FP6 IST consortium ‘EVERGROW’.

- [1] M. Mézard, G. Parisi, and M. A. Virasoro. *Spin-Glass Theory and Beyond*, volume 9 of *Lecture Notes in Physics*. World Scientific, Singapore, 1987.
- [2] R. G. Gallager. Low-density parity check codes. *IRE Trans. Inf. Theory*, IT-8:21, 1962.
- [3] D. J. C. MacKay. *Information theory, inference, and learning algorithms*. Cambridge University Press, Cambridge, 2003.
- [4] C. H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.
- [5] E. Friedgut. Sharp thresholds of graph properties, and the k -sat problem. *J. Amer. Math. Soc.*, 12, 1999.
- [6] M. Mézard, G. Parisi, and R. Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297:812–815, 2002.
- [7] M. Mézard and R. Zecchina. Random k -satisfiability problem: From an analytic solution to an efficient algorithm. *Phys. Rev. E*, 66:056126, 2002.
- [8] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina. Coloring random graphs. *Phys. Rev. Lett.*, 89:268701, 2002.
- [9] G. Semerjian and R. Monasson. A study of pure random walk on random satisfiability problems with “physical” methods. In E. Giunchiglia and A. Tachella, editors, *Proceedings of the SAT 2003 conference*, volume 120 of *Lecture Notes in Computer Science*, page 2919. Springer, 2004.

- [10] M. Mézard, T. Mora, and R. Zecchina. Clustering of solutions in the random satisfiability problem. *Physical Review Letters*, 94:197205, 2005.
- [11] T. Mora, M. Mézard, and R. Zecchina. Pairs of sat assignments and clustering in random boolean formulae, 2005. cond-mat/0506053.
- [12] D. Achlioptas and Y. Peres. The threshold for random k -sat is $2^k \log 2 - O(k)$. *Journal of the AMS*, 17:947–973, 2004.
- [13] R. Monasson. Optimization problems and replica symmetry breaking in finite connectivity spin-glasses. *J. Phys. A*, 31:515, 1998.
- [14] M. Mézard and G. Parisi. The bethe lattice spin glass revisited. *Eur. Phys. J. B*, 20:217, 2001.
- [15] F. Ricci-Tersenghi, M. Weigt, and R. Zecchina. Simplest random k -satifiability problem. *Phys. Rev. E*, 63:026702, 2001.
- [16] S. Cocco, O. Dubois, J. Mandler, and R. Monasson. Rigorous decimation-based construction of ground pure states for spin glass models on random lattices. *Phys. Rev. Lett.*, 90:047205, 2003.
- [17] O. Dubois and J. Mandler. The 3-xorsat threshold. *FOCS*, 00:769, 2002.
- [18] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina. Alternative solutions to diluted p -spin models and XORSAT problems. *J. Stat. Phys.*, 111:505, 2003.
- [19] T. Richardson and R. Urbanke. *Modern Coding Theory*. 2006. To be published, available at lthcwww.epfl.ch/mct.
- [20] H. Nishimori. *Statistical Physics of Spin Glasses and Information Processing: An Introduction*. Oxford University Press, Oxford, UK, 2001.
- [21] J. S. Yedidia, W. F. Freeman, and Y. Weiss. Constructing free energy approximations and generalized belief propagation algorithms. *technical report TR-2002-35, Mitsubishi Electrical Research Laboratories*, 2002. available at <http://www.merl.com>.
- [22] F. R. Kschischang, B. Frey, and H.-A. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Trans. Inform. Theory*, 47(2):498–519, 2001.
- [23] A. Montanari and G. Semerjian. On the dynamics of the glass transition on bethe lattices. cond-mat/0509366.
- [24] C. Di, A. Montanari, and R. Urbanke. Weight distributions of LDPC code ensembles: Combinatorics meets statistical physics. In *International Symposium on Information Theory*. IEEE, 2004.
- [25] C. Di, D. Proietti, I. E. Telatar, and R. L. Urbanke T. J. Richardson. Finite length analysis of low-density parity-check codes on the binary erasure channel. *IEEE Trans. Inform. Theory*, 48:1570–1579, 2002.
- [26] S. Condamin. Study of the weight enumerator function for a gallager code. 2002. <http://www.inference.phy.cam.ac.uk/condamin/report.ps>.
- [27] A. Vardy. The intractability of computing the minimum distance of a code. *IEEE Trans. Inf. Theory*, 43:1757–1766, 1997.
- [28] M. Mézard, M. Palassini, and O. Rivoire. Landscape of solutions in constraint satisfaction problems. *Phys. Rev. Lett.*, 95:200202, 2005.